

Fast Sparse Decomposition by Iterative Detection-Estimation

Arash Ali Amini,¹ *Student Member*,
Massoud Babaie-Zadeh,^{1*} *Member*,
and Christian Jutten¹ *Member*
(Last update: November 2006)

Abstract—Finding sparse solutions of underdetermined systems of linear equations is a fundamental problem in signal processing and statistics which has become a subject of interest in recent years. In general, these systems have infinitely many solutions. However, it may be shown that sufficiently sparse solutions may be identified uniquely. In other words, the corresponding linear transformation will be invertible if we restrict its domain to sufficiently sparse vectors. This property may be used, for example, to solve the underdetermined Blind Source Separation (BSS) problem, or to find sparse representation of a signal in an ‘overcomplete’ dictionary of primitive elements (i.e., the so-called atomic decomposition). The main drawback of current methods of finding sparse solutions is their computational complexity. In this paper, we will show that by detecting ‘active’ components of the (potential) solution, i.e., those components having a considerable value, a framework for fast solution of the problem may be devised. The idea leads to a family of algorithms, called ‘Iterative Detection-Estimation (IDE)’, which converge to the solution by successive detection and estimation of its active part. Comparing the performance of IDE(s) with one of the most successful method to date, which is based on Linear Programming (LP), an improvement in speed of about two to three orders of magnitude is observed.

I. INTRODUCTION

Finding (sufficiently) sparse solutions of underdetermined systems of linear equations has been studied extensively in recent years [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11]. The problem has a growing range of applications in signal processing. For example, it arises when dealing with underdetermined sparse source separation [7], [9], [11]. Another example is the so-called ‘atomic decomposition’ problem which aims at finding a sparse representation for a signal in an overcomplete dictionary [1], [2], [10]. Sparse representations are more suited for content analysis, i.e., extracting structure or meaning from a signal. They may also be used to achieve compression which in turn facilitates storage, processing and communication of signals. Recently, interesting applications have been reported in efficient (near-optimal) decoding of ‘error-correcting codes’ [12], [13], [14]. Also, some profound implications to the theory of sampling has been found [15], [16].

¹Electrical engineering department, Sharif university of technology, Tehran, Iran.

²Laboratoire des Images et des Signaux (LIS), Institut National Polytechnique de Grenoble (INPG), France.

This work has been partially funded by Sharif University of Technology, by French Embassy in Tehran, and by Center for International Research and Collaboration (ISMO).

Author’s email addresses are: aaamini57@yahoo.com, mbzadeh@yahoo.com and Christian.Jutten@inpg.fr

Corresponding author: Massoud Babaie-Zadeh, email: mbzadeh@yahoo.com, Tel: +98 21 66 16 59 25, Fax: +98 21 66 02 32 61.

It is not surprising that this fundamental problem has such a wide range of applications. The fact may simply be attributed to the widespread use of linear systems and transforms (throughout science and engineering). A linear transform (in a space with finite dimension) may be represented by a system of linear equations. Formerly, the underdetermined case, i.e., the case of ‘more unknowns than equations’, was considered degenerate and undesirable due to non-uniqueness of the solution. In other words, the corresponding linear transform is not invertible in this case which greatly reduces its usefulness for modeling (real-world problems). The general approach was (usually) to avoid the case by reformulating the underlying (physical) problem (to obtain enough equations in the unknowns). It is however possible to arrive at a unique solution by imposing additional constraints. One such constraint is (sufficient) sparsity of the solution, i.e., to require most components of the solution vector to be zero. More specifically, it can be shown that for a (random) system with n equations in $m(>n)$ unknowns, if there is a solution with less than $n/2$ (out of m) nonzero components, then it is (almost surely) the unique sparsest solution [17]. In other words, by limiting the domain of the underlying transform to ‘sufficiently sparse’ vectors, we can ensure its invertibility. We may even take one step further and claim that sparsity is usually more desirable than restrictive when it comes to signal processing applications. For example, in the context of atomic decomposition, a sparse solution leads to an efficient compact signal representation.

On the other hand, recent theoretical results [17] provide a solid mathematical basis for some of the methods (and optimality measures) experimentally found to produce sparse solutions. But issues still remain, perhaps one of the most important being the computational complexity of the available methods [7]. Our main objective in this paper is to introduce a framework which may be used to achieve fast sparse decomposition. But first, to get a better understanding of the problem, we review two contexts in which the problem arises, namely ‘Atomic Decomposition’ and ‘Sparse Component Analysis (SCA)’. Then we will review some of the available methods which will be used as a basis for comparison when evaluating the performance of our proposed method. We conclude the introduction with a brief layout of the rest of the paper.

In the atomic decomposition viewpoint [18], [1], we have ‘one’ signal whose samples are collected in the $n \times 1$ signal vector \mathbf{s} and the objective is to express it as a linear combination of a set of predetermined signals where their samples are collected in vectors $\{\phi_i\}_{i=1}^m$. After [18], the vectors ϕ_i are called *atoms* and the collection is called a *dictionary*. In mathematical language:

$$\mathbf{s} = \sum_{i=1}^m \alpha_i \phi_i = \Phi \boldsymbol{\alpha} \quad (1)$$

where Φ is the $n \times m$ dictionary matrix (with columns ϕ_i) and $\boldsymbol{\alpha}$ is the $m \times 1$ coefficient vector. To represent any $n \times 1$ vector, a basis of \mathbb{R}^n is sufficient, i.e., a collection of n linearly independent vectors (in \mathbb{R}^n). But if we take the number of atoms (much) more than it is required ($m \gg n$), then the likelihood that a given signal vector has a representation in terms of

only a few (i.e. much less than n) atoms is greatly increased. In that case, most of coefficients in the expansion would be negligible, i.e., the coefficient vector would be sparse. In fact with proper selection of dictionary, we may be able to find sparse representations for most of the signals of a signal space of interest. As mentioned before, such representations better reveal signal structure and are highly desirable from a practical point of view. A dictionary with $m > n$ atoms is called ‘overcomplete’ and the corresponding problem is usually referred to as ‘Atomic/Sparse Decomposition’ [1]. It is clear that this problem is essentially that of finding sparse solutions of an underdetermined linear system.

In the SCA viewpoint [7], [9], [19], [11], we use the sparsity assumption to solve the so-called ‘underdetermined’ Blind Source Separation (BSS) problem [20], [21]. The general BSS problem may be stated as: recovering m unknown source signals from n known mixtures of them, when little details are available about the sources and the mixing system. For example, usually the only assumption (or information) about the sources is their statistical independence. Similarly, regarding the mixing system only general properties (such as linearity/nonlinearity, being convolutive/instantaneous, ...) are assumed. Here, we only consider the most common mixing model, i.e. the (noiseless) linear instantaneous model:

$$\mathbf{x}(t) = \mathbf{A} \mathbf{s}(t), \quad t = 1, \dots, N$$

where $\mathbf{s}(t)$ and $\mathbf{x}(t)$ are the vectors containing sources and mixtures and \mathbf{A} is the (unknown) $n \times m$ mixing matrix. The only known quantity is $\mathbf{x}(t)$. The objective is to find the source vector and the mixing matrix only by observing $\mathbf{x}(t)$. For the case of ‘equal sources and mixtures’ ($m = n$) and with the assumption of an invertible mixing matrix \mathbf{A} , estimation of \mathbf{A} is sufficient to solve the problem. But in the underdetermined case where the number of sources is more than mixtures ($m > n$), even with the knowledge of \mathbf{A} , the system is not invertible and we are unable to obtain the sources. As mentioned before, this is where the added assumption of sparsity is helpful. More specifically, *if the original source vector \mathbf{s} is sufficiently sparse, then it is the unique sparsest solution of $\mathbf{x} = \mathbf{A}\mathbf{s}$* [17]. Again, the problem reduces to that of finding the sparse(st) solution of an underdetermined system. It is also interesting to note that ‘sparsity’ may also be used to estimate \mathbf{A} , by applying clustering techniques to the scatter plot of $\mathbf{x}(t)$ [22], [23]. We, however, assume \mathbf{A} to be known (or estimated) a priori. Moreover, we assume that the energy of the columns of \mathbf{A} are normalized to 1, that is $\|\mathbf{a}_i\|_2 = \mathbf{a}_i^T \mathbf{a}_i = 1$ (this is always possible because as it is seen in (1), each ϕ may be multiplied by a scalar and the corresponding coefficient divided by that scalar. In BSS, this is usually called “scale indeterminacy”). It is also notable that sparsity is not much of a restriction in practice: Many natural signals exhibit sparsity either in the time-domain or in a transform-domain [19], [7], [9].

For future discussions, we will mainly adopt the terminology and notation of SCA, although some references might be made to the atomic decomposition terminology. This is partly because nearly all the methods to be reviewed have been originally developed in the context of atomic decomposition.

The methods used for sparse decomposition may be divided into two categories: those selecting a solution of the underdetermined system by *minimizing a cost function* over the space of all possible solutions, and those taking a more *algorithmic approach* without explicitly specifying a cost function. For the methods of the first type, the cost function may be viewed as a *measure of sparsity*¹ of the solution vector. One such measure, which is strongly supported by our intuition of sparsity (and may even be considered its definition), is the so-called l^0 norm of \mathbf{s} denoted by $\|\mathbf{s}\|_0$ and defined as the number of nonzero elements of \mathbf{s} . Unfortunately, minimizing the l^0 norm requires combinatorial search which quickly becomes intractable as the dimension increases; It is also highly sensitive to noise. It has been shown first experimentally [1] and then theoretically [2], [17], [3], [4], [5], [6], [8] that the l^0 norm could be replaced by l^1 norm, i.e., we seek a solution minimizing $\|\mathbf{s}\|_1 = \sum_{i=1}^m |s_i|$. The l^1 norm is more robust to noise and more importantly, the associated optimization problem is ‘convex’ which can be solved much more efficiently. The problem may also be stated as a Linear Programming (LP) problem and then solved in polynomial time using interior-point methods. Minimizing l^1 norm, which was initially named Basis Pursuit (BP), may be considered the most successful method to date. We will refer to this method as the ‘LP approach’ to emphasize that we will use linear programming techniques (mostly interior-point solvers) to obtain its solution.

Besides LP, we also consider two other earlier approaches to atomic decomposition. One of them, which we denote as the Method of Frames (MOF) following [1], obtains a solution of $\mathbf{x} = \mathbf{A}\mathbf{s}$ having minimal l^2 norm, i.e., $\|\mathbf{s}\|_2 = (\sum_{i=1}^m s_i^2)^{1/2}$. The method has been originally developed without any regard of sparsity [24], and it turns out that its solution is usually not sparse. But merely as a method of decomposition, it has some nice properties: the solution is linear in \mathbf{x} and it may be obtained using the pseudo-inverse of \mathbf{A} , i.e., $\mathbf{s}_{\text{MOF}} = \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{x}$. It may also be considered as the best linear inverse system in the Least Squares (LS) sense (both statistically and deterministically). We will mainly use MOF as a benchmark for the speed of algorithms².

The other approach is Matching Pursuit (MP) developed by Mallat and Zhang [18] (who also coined the name atomic decomposition). It may be considered an algorithmic approach and one of the first methods to target sparsity of the solution (though implicitly). Recall that in the atomic decomposition we seek a linear expansion of \mathbf{x} in terms of atoms ϕ_i . MP begins by finding the best single-atom approximation of \mathbf{x} in the LS sense, i.e., $\mathbf{x} \approx \hat{\mathbf{x}}_1 = s_1 \phi_1$ where s_1 and ϕ_1 are selected such that $\|\mathbf{x} - s_1 \phi_1\|_2$ is minimized (over $1 \leq i \leq m$). This is equivalent to finding the atom which best correlates with \mathbf{x} , i.e., for which $|\mathbf{x}^T \phi_i|$ is maximum. If the residue $\mathbf{x} - \hat{\mathbf{x}}_1$ is small enough, the algorithm is terminated, otherwise the

¹To be more precise, the cost function should be considered as a measure of deviation (or departure) from sparsity, but for the sake of simplicity we will neglect such formality.

²Because of the existence of highly efficient numerical algorithms for the computation of pseudo-inverse (with computational cost close to solving a linear system of comparable dimensions), MOF may be considered to achieve fastest decomposition.

same process is repeated for the residue. In other words, at each step, MP finds the best single-atom approximation of the residue. In this sense, it is a *greedy algorithm* (selecting the best choice given the current situation). We have a good chance of obtaining a sparse representation if the algorithm terminates early (i.e., with a number of atoms much less than m). However, as with any greedy algorithm, there are situations in which an early mistake would lead to large deviation from the optimal solution. We will discuss this issue further in the experimental results section.

Among methods of decomposition available, the fast methods (e.g. MP or MOF) usually don't produce accurate results, while LP which is guaranteed to obtain the exact solution (asymptotically) will become very computationally demanding at large dimensions. Our proposed algorithm (or framework) is an attempt to keep accuracy while approaching MP and MOF in speed. We begin with a general introduction of the 'Iterative Detection-Estimation (IDE)' framework, followed by a detection-theoretic motivation for the derivation of IDE algorithms. We then develop two versions of such algorithms denoted as 'IDE-s' and 'IDE-x', followed by some comments on the choice of parameters. We conclude with a discussion of experimental results comparing the performance of the proposed algorithms to existing methods.

II. ITERATIVE DETECTION-ESTIMATION

Perhaps one of the main obstacles to implementation of many optimal methods of sparse decomposition is the inherent 'combinatorial search' required. The obstacle is overcome if we could somehow detect which components of the (original) source vector \mathbf{s} are 'active'. By active sources we mean those having a *considerable value*, as opposed to those being *nearly zero* and denoted as being 'inactive'. The key idea here is to detect (or determine) the 'activity' status of each source separately (i.e., independently of all the other sources). The total number of detections required would be m which is linear in the problem dimension.

The problem with this approach is that optimal detection of 'activity' of a source requires exact knowledge of the values of other sources. Our solution is to use a suboptimal detector with the exact values replaced by some previously known estimate (or an initial guess). This rough detection may (surprisingly) be used to obtain a better estimate of source vector which in turn may be used to enhance the detection. By iteratively applying a detection-step followed by an estimation-step³ we can hopefully get progressively better estimates and get closer to the original source vector, hence the name 'Iterative Detection-Estimation (IDE)'. This convergence will be justified by our experimental results, although the theoretical convergence proof is a tricky and open question.

Fig. 1 illustrates a schematic diagram of the algorithm in its general form. In this figure, k is the iteration index, $\hat{\mathbf{s}}^{(k)}$ and $\hat{\mathbf{s}}^{(k+1)}$ are respectively current and next estimate of the source vector, and \mathcal{I}_α is the set of indices of the sources detected to

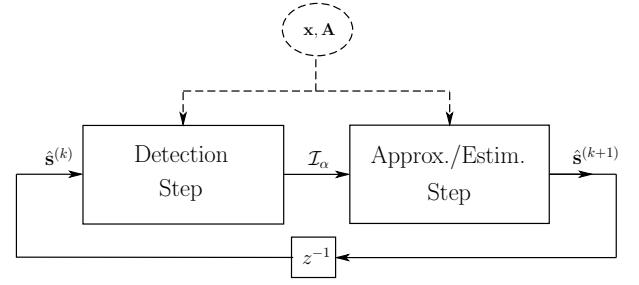


Fig. 1. Schematic diagram illustrating IDE operation: $\hat{\mathbf{s}}^{(k)}$ is the source vector estimate after k -th iteration; \mathcal{I}_α denotes the (set of) indices of sources detected active.

be active⁴.

We begin the discussion by giving a motivation for the detection step based on a simple (statistical) model of sparsity. We then give two versions of the estimation step leading to two versions of IDE, namely IDE-s and IDE-x. Throughout the discussion, k_α and k_ι will be used to denote the number of sources detected active and inactive, respectively. Also, throughout the development of the algorithm, the term 'active (inactive) sources' usually means those sources 'detected active (inactive)'. We sometimes use it to refer to original 'active (inactive) source'. The distinction should be apparent from the context.

III. DETECTION STEP

A. motivation

To provide motivation for the detection step, we first use a Mixtures of Gaussians (MoG) to distinguish active/inactive states of a sparse source. This provides us with a simple (intuitive) model of sparsity. More specifically, let π_0 be the probability of s_i being *inactive* ($\pi_0 \lesssim 1$ to insure sparsity). Then, the value of an inactive source is modeled by $\mathcal{N}(0, \sigma_0^2)$, and an active source by $\mathcal{N}(0, \sigma_1^2)$, where $\sigma_0^2 \ll \sigma_1^2$ ⁵. The probability π_0 will not be used in the development of the algorithm, but will be useful as a measure of sparsity in the experimental results.

As stated previously, we detect the activity status of each source separately. Assume that we want to determine the status of the i -th source s_i . We observe $\mathbf{x} = s_i \mathbf{a}_i + \sum_{j \neq i} s_j \mathbf{a}_j$ and we wish to decide which of the following two hypotheses has occurred :

$$\begin{aligned} H_0 : & \quad s_i \sim \mathcal{N}(0, \sigma_0^2), \\ H_1 : & \quad s_i \sim \mathcal{N}(0, \sigma_1^2). \end{aligned}$$

This is essentially a binary hypothesis testing problem [25]. It may be argued that $t_i = \mathbf{a}_i^T \mathbf{x}$ contains all the information regarding the discrimination of the two hypothesis, i.e., it is a sufficient statistic for the problem (given the value of all the other sources). Defining $\mu_i \triangleq \sum_{j \neq i} s_j \mathbf{a}_i^T \mathbf{a}_j$ and noting that $t_i = s_i + \mu_i$, we can reformulate the problem in terms of the sufficient statistic as $H_k : t_i \sim \mathcal{N}(\mu, \sigma_k^2)$ for $k = 0, 1$.

⁴Subscript α is used to designate quantities related to active sources. Similarly, subscript ι is used for inactive sources.

⁵A shorthand notation would be $s_i \sim \pi_0 \mathcal{N}(0, \sigma_0^2) + (1 - \pi_0) \mathcal{N}(0, \sigma_1^2)$

³This step may also be called approximation or projection step depending on the approach we use to obtain the estimate.

We approach the problem in the Neyman-Pearson framework, considering $\{s_j\}_{j \neq i}$ to be parameters (rather than random variables). Also, we do not assign priors to the hypotheses. The optimal test (in the NP sense) would then be a likelihood ratio test, i.e., one which compares the likelihood ratio to a threshold. For the problem at hand the critical region of this test may be written as

$$\log \frac{\sigma_0}{\sigma_1} + \left(\frac{1}{2\sigma_0^2} - \frac{1}{2\sigma_1^2} \right) (t_i - \mu_i)^2 > \tau$$

or after absorbing known constants into the threshold as

$$|t_i - \mu_i| > \epsilon$$

where ϵ is the new threshold. Recalling the definition of μ_i , it is observed that implementing the optimal test for activity of s_i requires knowledge of all the other sources⁶. As mentioned before, our solution is to replace them with their estimates (obtained from a previous iteration or from an initial guess). The resulting sub-optimal test is then

$$g_i(\mathbf{x}, \hat{\mathbf{s}}) \triangleq \left| \mathbf{a}_i^T \mathbf{x} - \sum_{j \neq i} \hat{s}_j \mathbf{a}_j^T \mathbf{a}_j \right| > \epsilon$$

for $1 \leq i \leq m$. We will call $g_i(\mathbf{x}, \hat{\mathbf{s}})$ as defined above the activity function associated with the i -th source. Below, we have summarized the detection step where we have also allowed the threshold to vary with iteration. It is found experimentally that decreasing the threshold each iteration produces better results.

Detection Step : Obtain active indices according to

$$\mathcal{I}_\alpha = \{1 \leq i \leq m : g_i(\mathbf{x}, \hat{\mathbf{s}}^{(k)}) > \epsilon^{(k+1)}\}$$

B. vector form

It is possible to write the detection step in a simple form using vector-matrix notations. Note that one may write the activity function as

$$\begin{aligned} g_i(\mathbf{x}, \hat{\mathbf{s}}) &= |\mathbf{a}_i^T (\mathbf{x} - \mathbf{A}\hat{\mathbf{s}} + \mathbf{a}_i \hat{s}_i)| \\ &= |\mathbf{a}_i^T (\mathbf{x} - \mathbf{A}\hat{\mathbf{s}}) + \hat{s}_i|. \end{aligned}$$

If one collects the components $g_i(\mathbf{x}, \hat{\mathbf{s}})$ in a ‘vector activity function $\mathbf{g}(\mathbf{x}, \hat{\mathbf{s}})$ ’, the detection step may simply be stated as

$$\mathbf{g}(\mathbf{x}, \hat{\mathbf{s}}) = |\mathbf{A}^T (\mathbf{x} - \mathbf{A}\hat{\mathbf{s}}) + \hat{\mathbf{s}}| > \epsilon$$

where $|\cdot|$ and $>$ operate component-wise when used on vectors. Note that if the previous estimate $\hat{\mathbf{s}}$ is itself a solution of the system (i.e. $\mathbf{x} = \mathbf{A}\hat{\mathbf{s}}$), then the (vector) activity function is simply $\mathbf{g}(\mathbf{x}, \hat{\mathbf{s}}) = |\hat{\mathbf{s}}|$ (this is the case for IDE-s algorithm discussed below). But the previous estimate does not need to satisfy the system, in which case the term $\mathbf{A}^T (\mathbf{x} - \mathbf{A}\hat{\mathbf{s}})$ acts as a compensator (this is the case for IDE-x). Also note that (as a special case) the activity function evaluated at the true source vector is $\mathbf{g}(\mathbf{x}, \mathbf{s}) = |\mathbf{s}|$. This is useful when selecting threshold values.

⁶ Note that because of the dependence of the critical region on the value of μ_i , there is no Uniformly Most Powerful (UMP) test.

IV. ESTIMATION STEP

Knowing the sparsity pattern (i.e. active index set \mathcal{I}_α), the estimation of sources would be straightforward. Here, we introduce two simple approaches which may be considered respectively as projections in the source space (s-space) and the mixture space (x-space).

A. s-space approach

In this approach we obtain the source vector by solving the following optimization problem:

$$\hat{\mathbf{s}} = \arg \min_{\mathbf{s}} \sum_{i \in \mathcal{I}_l} s_i^2 \quad (\text{s.t. } \mathbf{x} = \mathbf{A}\mathbf{s}) \quad (2)$$

where $\mathcal{I}_l = \mathcal{I}_\alpha^c$ is the inactive index set. Let $k_\alpha \triangleq |\mathcal{I}_\alpha|$ ($k_l \triangleq |\mathcal{I}_l| = m - k_\alpha$) be the number of sources detected active (inactive). The above operation may be thought of as projection into the (k_α -dimensional) subspace determined by the active indices. We denote the IDE algorithm using this approach for source estimation as ‘IDE-s’.

Optimization problem (2) is a special case of Quadratic Programming (QP) which has been extensively studied in the literature [26]. For simplicity, assume (for the rest of this section) that ‘the first k_l sources’ have been detected inactive, i.e., $\mathcal{I}_l = \{1, 2, \dots, k_l\}$. Then, the cost function in (2) may be stated as the quadratic form $\mathbf{s}^T \mathbf{H} \mathbf{s}$ with $\mathbf{H} = \begin{pmatrix} \mathbf{I}_{k_l} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$ where \mathbf{I}_{k_l} is the $k_l \times k_l$ identity matrix.

Among the many numerically efficient approaches available [26], [27], here we consider direct solution of the so-called Karush-Kuhn-Tucker (KKT) system of equations which serves as a necessary condition for optimality [26], i.e., the optimal solution should satisfy

$$\begin{pmatrix} \mathbf{H} & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{s} \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{x} \end{pmatrix}$$

where $\boldsymbol{\lambda}$ is the $n \times 1$ vector of Lagrange multipliers. Under certain conditions, explicit formulas for the solution of this system may be obtained. Partitioning vectors and matrices into ‘inactive/active’ parts, we have

$$\begin{pmatrix} \mathbf{I}_{k_l} & \mathbf{0} & \mathbf{A}_l^T \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_\alpha^T \\ \mathbf{A}_l & \mathbf{A}_\alpha & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{s}_l \\ \mathbf{s}_\alpha \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{x} \end{pmatrix}.$$

Under the fairly general condition of $k_\alpha \leq n$, the ‘unique’ solution of the problem may be stated as

$$\begin{cases} \hat{\mathbf{s}}_l = \mathbf{B}_l^T (\mathbf{B}_l \mathbf{B}_l^T)^{-1} \mathbf{Z}^T \mathbf{x} \\ \hat{\mathbf{s}}_\alpha = (\mathbf{A}_\alpha^T \mathbf{A}_\alpha)^{-1} \mathbf{A}_\alpha^T (\mathbf{x} - \mathbf{A}_l \mathbf{s}_l) \end{cases} \quad (3)$$

where \mathbf{Z} is a $n \times (n - k_\alpha)$ matrix whose columns form a basis for the null space of \mathbf{A}_α^T , and $\mathbf{B}_l \triangleq \mathbf{Z}^T \mathbf{A}_l$. Another closed-form solution may be obtained under a more restrictive condition, namely $k_\alpha \leq \min\{n, m - n\}$. It can be shown [28] that in this case, the ‘unique’ solution of the problem is

$$\begin{cases} \hat{\mathbf{s}}_\alpha = (\mathbf{A}_\alpha^T \mathbf{P} \mathbf{A}_\alpha)^{-1} \mathbf{A}_\alpha^T \mathbf{P} \mathbf{x} \\ \hat{\mathbf{s}}_l = \mathbf{A}_l^T \mathbf{P} (\mathbf{x} - \mathbf{A}_\alpha \mathbf{s}_\alpha) \end{cases} \quad (4)$$

where $\mathbf{P} \triangleq (\mathbf{A}_l \mathbf{A}_l^T)^{-1}$. Obtaining the solution using these two explicit formulas is usually faster than directly solving

the $(m+n) \times (m+n)$ KKT system. In the experiments of this paper, the second closed-form (4) will be used.

B. x -space approach

In this approach, the source estimate is obtained as the solution of the following optimization problem:

$$\begin{cases} \hat{\mathbf{s}}_\alpha = \arg \min_{\mathbf{s}_\alpha} \|\mathbf{x} - \mathbf{A}_\alpha \mathbf{s}_\alpha\|_2 \\ \hat{\mathbf{s}}_l = \mathbf{0} \end{cases} \quad (5)$$

In other words, we estimate the active part of the source vector by projecting \mathbf{x} into the subspace spanned by the (allegedly) active atoms, and simply set the inactive part to zero. Since this expansion of \mathbf{x} in terms of the active atoms occurs in the mixture space, we denote the associated IDE method as ‘IDE-x’.

Using pseudo-inverse of \mathbf{A}_α , the solution of (5) may simply be stated as (assuming $k_\alpha \leq n$)

$$\begin{cases} \hat{\mathbf{s}}_\alpha = (\mathbf{A}_\alpha^T \mathbf{A}_\alpha)^{-1} \mathbf{A}_\alpha^T \mathbf{x} \\ \hat{\mathbf{s}}_l = \mathbf{0} \end{cases} \quad (6)$$

It is interesting to note that setting $\hat{\mathbf{s}}_l$ to zero in (3) also leads to the same result. Since we only care about true values of active sources and expect inactive ones to be nearly zero, this is a reasonable simplification. In this sense, IDE-x may be considered an approximation of IDE-s. It is important to note that the IDE-x solution no longer satisfies $\mathbf{x} = \mathbf{A}\mathbf{s}$, and hence as later experiments show, this slightly lowers the accuracy of IDE-x relative to IDE-s. The loss is, however, negligible when the noise over the inactive part of the solution is not high. This is the price we pay for the tremendous gain in speed obtained due to the simplified structure of the IDE-x estimate.

V. INITIAL CONDITIONS

To initiate iterations, we need an initial estimate. For all the experiments in this paper, we will use the simplest initial condition, i.e. $\hat{\mathbf{s}}^{(0)} = \mathbf{0}$. Note that this is not a solution of $\mathbf{x} = \mathbf{A}\mathbf{s}$, but as was mentioned before, the detection step does not require the (initial or middle) estimates to satisfy the system. Also note that in the absence of prior information, the ‘zero initial condition’ is perhaps the most reasonable one, because due to the sparse nature of the actual solution, most sources would be zero anyway.

One may also use other ‘cheap’ estimates initially. For example, IDE may be used to improve upon the solution of the MOF method.

VI. COMMENTS ON THE CHOICE OF THRESHOLDS

In this section, we briefly discuss some issues regarding threshold selection. First consider the ideal case where the ‘actual’ inactive sources are (exactly) zero. Now suppose that 1) the detection step ‘at least’ detects the actual active sources correctly (there might also be some actually inactive ones, incorrectly detected active). Then if 2) the solution of the estimation step is unique, it will coincide with the actual sparse solution since the latter achieves a cost function value of zero (for both IDE-s and IDE-x). In other words, the estimation

step compensates for the mistakes made during detection and correctly estimates ‘all’ inactive sources to be zero. One way to guarantee the uniqueness of the solution [for either of (2) or (5)] is by keeping the number of sources detected active below the number of mixtures (i.e. $k_\alpha < n$).

The two conditions above suggest that there are *implicit bounds* on the value of threshold. It should be low enough to guarantee that (nearly) all the actual active sources are detected correctly. On the other hand, it should be high enough to keep the number of those detected active below n . The above argument then suggests that within those bounds a rough detection is sufficient and will lead to the desired solution. In practice, for the moderately difficult problem⁷ those bounds provide enough gap for us to easily select thresholds. As will be seen in the experimental section, it may even be possible to obtain threshold sequences which work well for ‘families’ of problems. We will also see that IDE is even robust to errors in detection of actual active sources, in the sense that minor ‘missed detections’ are corrected through iteration.

There are also *explicit bounds* on the threshold. Recall that $g_i(\mathbf{x}, \mathbf{s}) = |s_i|$. This suggests that any bound on the the absolute value of the sources would translate (somewhat directly) into a bound on the threshold. One might then restrict the threshold to $0 < \epsilon < K \cdot \|\mathbf{s}\|_\infty$ where $K \gtrsim 1$ (values of K greater than unity may be used to account for estimation errors). For simplicity, in all the experiments of this paper, we will assume that the original source vector is normalized to unit l^∞ norm (i.e. $\|\mathbf{s}\|_\infty = 1$) and then select thresholds in the interval $(0, 1)$ (i.e. $K = 1$). In real applications, one needs to estimate $\|\mathbf{s}\|_\infty$. One simple approach is to take the activity function at the first iteration as an estimate of source absolute value. Thus if the ‘zero initial condition’ is used one gets the estimate $\|\mathbf{g}(\mathbf{x}, \hat{\mathbf{s}}^{(0)})\|_\infty = \|\mathbf{A}^T \mathbf{x}\|_\infty$.

VII. EXPERIMENTAL RESULTS

In this section, we will examine the performance of the two versions of IDE, i.e., IDE-s and IDE-x, and compare them to some of the available methods. This will be done by discussing the results of five experiments detailing different aspects of IDE behavior.

In all the experiments, the \mathbf{A} matrix will be generated randomly by drawing each of its m columns from a uniform distribution on the unit sphere in \mathbb{R}^n . We will use the Gaussian mixture model discussed earlier to generate source vectors in the first three experiments. A different source model will be used for the last two experiments which will be explained later. In any case, we always normalize the source vector so that $\|\mathbf{s}\|_\infty = 1$. This limits the choice of thresholds to the interval $(0, 1)$.

We will use SNR as a measure of quality (or accuracy) of the solution produced by an algorithm. To measure complexity, the total CPU time required by the algorithm will be used (although this is not an exact measure of complexity, it provides us a rough estimation). Depending on the context, two different forms of SNR will be considered. When dealing

⁷A sparse decomposition problem gets difficult when n/m decreases or the actual solution becomes less sparse.

TABLE I
IDE PROGRESS TOWARD FINAL SOLUTION

k	$\epsilon^{(k)}$	IDE-s			IDE-x		
		k_α	ΔT	SNR	k_α	ΔT	SNR
1	0.3	158	0.377	6.44	158	0.025	5.50
2	0.2	47	0.297	8.24	49	0.008	8.24
3	0.1	58	0.292	11.85	149	0.019	14.51
4	0.05	73	0.293	18.26	96	0.013	21.06
5	0.02	105	0.310	25.36	176	0.026	27.88
6	0.01	107	0.315	30.27	126	0.021	28.80

with a single realization (or sample) of the system $\mathbf{x} = \mathbf{A}\mathbf{s}$, we usually use what may be called ‘Spatial SNR (SSNR)’, which is defined as $\|\mathbf{s}\|_2^2 / \|\mathbf{s} - \hat{\mathbf{s}}\|_2^2$ where \mathbf{s} and $\hat{\mathbf{s}}$ are respectively the original and the estimated source vectors⁸. Since we are dealing mostly with large systems (e.g. $m = 500, 1000$), this form of averaging is justified. When working with many samples of the system $\{\mathbf{x}(t) = \mathbf{A}\mathbf{s}(t)\}_{t=1}^N$, we usually average over time (index) obtaining ‘Temporal SNR’ for each source, i.e.,

$$(\text{Temporal}) \text{SNR}_i = \frac{\sum_{t=1}^N s_i^2(t)}{\sum_{t=1}^N [s_i(t) - \hat{s}_i(t)]^2}, \quad 1 \leq i \leq m.$$

The context indicates which SNR definition is being used, and hence, we often omit the ‘spatial’ or ‘temporal’ prefixes.

For the purpose of comparison, three of the available decomposition methods, namely MOF, MP, and LP, will be considered. The emphasis is on LP since this is the one guaranteed to obtain the spars(est) solution. In all the experiments, unless explicitly stated otherwise, the LP solution is obtained using MATLAB 7.0 implementation of an ‘interior-point’ LP solver (called LIPSOL). Also, all the CPU times are measured on a 2.4GHz P4 CPU under MATLAB 7.0 environment.

A. experiment 1 - evolution toward the solution

1) *a typical setting*: In this experiment, we will study the typical behavior of IDE by considering a ‘single’ realization of a system with dimensions $m = 1024$ and $n = \lfloor 0.4m \rfloor = 409$. The source vector is drawn from a Gaussian mixture with $\pi_0 = 0.9$, $\sigma_0/\sigma_1 = 0.01$ and is normalized so that $\|\mathbf{s}\|_\infty = 1$. In a single realization, the actual number of active components in the source vector is more important than the π_0 parameter (which somehow measures sparsity ‘on the average’). In particular, for the (random) source vector considered here, the number of sources with absolute values over 0.01 is obtained to be 105. This is nearly equal to $n/4$ which signifies a relatively difficult problem (as will be proposed by experiment 4).

Both versions of the IDE algorithm have been applied to the problem. In either case, a total number of six iterations has been used with threshold values $\epsilon = 0.3, 0.2, 0.1, 0.05, 0.02, 0.01$. This sequence has been found experimentally to produce results as accurate as those of LP, for the problem family characterized by ($\pi_0 = 0.9, n/m = 0.4$).

⁸Note that here we average over the source (or spatial) index, on a single time sample.

TABLE II
COMPARISON OF DIFFERENT ALGORITHMS

algorithm	total CPU time	SNR (dB)
IDE-s (6 itrs.)	$1.88 e 00$	30.27
IDE-x (6 itrs.)	$1.12 e -1$	28.80
LP (interior-pt)	$1.23 e +2$	26.25
LP (simplex)	$5.45 e +3$	26.25
MP (10 itrs.)	$1.54 e -1$	1.80
MP (100 itrs.)	$1.58 e 00$	10.70
MP (1000 itrs.)	$8.71 e 00$	9.82
MOF	$1.38 e -1$	2.36

The results obtained at the end of each iteration are summarized in Table I. For each of the IDE-s and IDE-x, the number of sources detected active (k_α), the elapsed CPU time in seconds, and the (spatial) SNR, all obtained at the end of each iteration have been recorded. Also, Fig. 2 provides a more visual account of IDE-s progress toward the solution (the progress of IDE-x is similar). Each plot in this figure shows the original and the estimated source vectors after an iteration, respectively designated by small black and large gray dots. The vectors are plotted against the source index (i.e., the plots are s_i or \hat{s}_i versus i). We have also identified sources detected to be active after each iteration by drawing a small square above them.

Based on these results we can make the following observations: At first, due to the low starting threshold value ($= 0.3$), the number of sources detected active is more than necessary ($158 = k_\alpha > \text{actual \# act.} \approx 105$). The number, however, satisfies the uniqueness condition (of the estimation part) $k_\alpha < n$ which enables IDE(s) to start the iteration. Also note from the figure that (for IDE-s) not all the actual active sources are at first detected. The figure shows that *the initial guess for active sources is highly improved after the second iteration* and this improvement continues (though more gradually) until the algorithm converges to the original solution. There are also (a few) sources correctly detected active at first, wrongly discarded at a later iteration, but eventually re-detected at final iterations. This shows the self-correcting capability of IDE; A property that a greedy algorithm such as MP does not possess.

Note that for IDE-s, the final number of sources detected active is near the actual value. For IDE-x, final k_α is higher, but the final solution has the same quality ($\text{SNR} \approx 10^3$ or 30 dB). This is in accordance with our previous statement that false alarm in detection of active sources does not affect the performance as long as it remains within the limits of the uniqueness condition.

Another notable observation is that for both versions, SNR increases by nearly an order of magnitude every two iterations until it reaches the final value of $\approx 10^3$ which as we will see is comparable to the quality obtainable by LP. Also note that *each iteration of IDE-x is nearly an order of magnitude faster than that of the IDE-s*; A property that holds in general as will be confirmed in a later experiment.

2) *comparison of algorithms*: In Table II, we have summarized the results obtained by some of the available methods when applied to the same realization of the problem (along with those of IDE’s). For LP, both the interior-point and Simplex implementations are considered. For MP, the results




Fig. 2. Progress of IDE-s toward final solution (Experiment 1) : $m = 1024$, $n = \lfloor 0.4m \rfloor = 409$, $\#act \approx 105$. Each plot shows the original source vector (black) and its estimate obtained after an iteration (gray). The sources detected to be active are marked with a black square above each plot. Six iterations were used with threshold values (from top to bottom) $\epsilon = 0.3, 0.2, 0.1, 0.05, 0.02, 0.01$. The top plot corresponds to the first iteration. -

after 10, 100, and 1000 iterations are recorded separately.

It is observed that both versions of IDE achieve a final SNR of nearly 30 dB (after six iterations) which is slightly better than the 26 dB obtained by LP. The major difference is in the time required by each algorithm. In fact, with nearly the same final SNR, the time comparison would be more meaningful.

We observe that IDE-x is ten times faster than IDE-s which itself is a hundred times faster than LP-interior which in turn is ten times faster than LP-Simplex. Thus, IDE-x, for example, achieves nearly four orders of magnitude improvement in speed over LP-Simplex, which is a truly remarkable achievement. The average results are more or less the same, as will be discussed in the third experiment.

A comparison with the results obtained by MOF shows that it has nearly the same speed as IDE-x. The final quality achieved (≈ 2 dB) is however far from acceptable. This is not surprising since MOF was not meant originally to select the spars(est) solution.

The quality and time obtained by MP after 10 iterations is very close to those of MOF. The best performance is achieved around 100 iterations with a final SNR value of nearly 11 dB and a time comparable to that of IDE-s. This is the maximum quality attainable by MP. It may partly be explained by recalling that in the present problem, the number of (actual) active sources is nearly 105 and that for MP, the number of (active) atoms present in the expansion (of \mathbf{x}) is the same as the number of iterations. The claim is further confirmed by noting that after 1000 iterations the quality actually degrades

to ≈ 10 dB. The observation reveals the fundamental problem of ‘greedy algorithms’ of which MP is one. We will discuss the problem shortly and show how IDE-x effectively evades it.

3) *IDE-x versus MP*: Before concluding this experiment, we want to briefly comment on how IDE-x may be used to improve upon MP. There is a resemblance between the two algorithms. Recall that, at each step, MP finds the atom that best correlates with the residue (up to that point). In this sense, MP finds successive ‘single-atom approximations’ to \mathbf{x} which at the end add up to be build the final estimate. In contrast, at each iteration, IDE-x expands \mathbf{x} over all the atoms detected to be active, and hence, it is more likely to obtain the optimal (sparse) expansion.

Fig 3 shows that this is indeed the case. In this figure, the relative approximation error in the expansion of \mathbf{x} is plotted versus iteration (or step) for both IDE-x and MP. Note that MP requires nearly 1000 steps to achieve the same error that IDE-x has achieved in 6 iterations. Moreover, in doing so, MP incorporates into the expansion nearly all the 1024 atoms available (recall that for MP each step adds one atom). Consequently, the resulting \mathbf{s} vector is far from sparse. This reflects the main problem of greedy algorithms: making an early mistake usually takes many steps to correct, during which the algorithm deviates considerably from the optimal solution. IDE-x (and in general IDE’s) avoid this by expanding over all possible candidates at each iteration.




Fig. 3. IDE-x versus MP: Relative approximation error in \mathbf{x} obtained by IDE-x/MP at each iteration/step plotted against iteration/step index (k). The data is from experiment 1.




Fig. 4. Temporal SNR versus the source index ($1 \leq i \leq m$) for the three algorithms IDE-s, IDE-x and LP, in three different settings: (top) $(m, n/m) = (100, 0.6)$, (middle) $(m, n/m) = (500, 0.6)$ and (bottom) $(m, n/m) = (500, 0.4)$. Temporal averages are over $N = 1000$ samples. All samples are drawn from a Gaussian mixture model (for the sources) with $\pi_0 = 0.9$ and $\sigma_0/\sigma_1 = 0.01$.

B. experiment 2 - average quality

In this experiment, we compare average behavior of IDE's with that of LP. The three algorithms are applied to $N = 1000$ time samples $\{\mathbf{x}(t)\}_{t=1}^N = \{\mathbf{A}\mathbf{s}(t)\}_{t=1}^N$. The 'temporal SNR' is then obtained for each algorithm and plotted against the source index (i.e. (Temporal) SNR_i versus i). Fig. 4 shows the results for three illustrative cases. For all the cases a Gaussian mixture model with $\pi_0 = 0.9$, $\sigma_0/\sigma_1 = 0.01$ is used to generate the N time samples. The three plots correspond to different choices of $(m, n/m)$ pairs, i.e. $(100, 0.6)$, $(500, 0.6)$ and $(500, 0.4)$ respectively.

A fixed threshold sequence, namely $\epsilon = 0.7, 0.6, 0.5, 0.4, 0.3, 0.2, 0.1, 0.07, 0.05, 0.02$, is used in all the three cases and over all the N samples. This sequence is found (experimentally) to produce slightly better results than LP in all cases of interest. *Note that although we have set the*




Fig. 5. Average CPU time in sec. versus problem dimension (m) for various algorithms. At all dimensions, $n = 0.4m$. Temporal averages are over $N = 10$ samples.

thresholds manually, they are only set once at the beginning and there is no need to change them on a per-sample basis. Also more experiments with other combinations of the problem parameters (i.e. $(m, n/m, \pi_0)$) showed that this is indeed a 'good' choice for nearly all problems for which LP is 'good', especially at higher dimensions (i.e. for large m).

The three cases in Fig. 4 were chosen to illustrate some general trends. Note that IDE's outperform LP as shown by the gap between their average (temporal) SNRs, but the gap reduces as the dimension is increased (i.e. increasing m while n/m is fixed). In other words, the performance of the algorithms converges to one another as we increase m . This is confirmed by more experiments. Another trend is that the gap is usually reduced as the problem gets harder (i.e. decreasing n/m while m is fixed). The third plot also shows that surprisingly sometimes IDE-x (slightly) outperforms IDE-s.

C. experiment 3 - average complexity

In this experiment, we will examine the relative complexity (or speed) of the algorithms more closely. The measure to be used is the 'average CPU time' required by each algorithm. More specifically, we are interested in 'average time' versus 'problem dimension' plots where the dimension is m , the number of sources. We select seven points in the interval⁹ $10 \leq m \leq 10^3$, and for each m , we generate $N = 10$ instances of the problem, keeping n/m fixed at nearly 0.6 (or more exactly $n = \lfloor 0.6m \rfloor$). Each of the algorithms under study is then applied to the N samples and the average time (obtained over the N samples) is used as an index of complexity at the specified dimension. Fig. 5 summarizes the results.

To generate the figure, all the iterative algorithms (i.e., IDE-s, IDE-x and MP) have been applied only for 10 iterations. Moreover, we have only considered the interior-point implementation of LP.

Examining the figure, similar patterns as those encountered earlier may be identified. Again, the slowest algorithm is LP followed by IDP-s which is more than one order of magnitude faster; The difference being nearly constant across dimension. It is interesting to note that IDE-x may be grouped along with

⁹The points are selected to be equidistant in the logarithmic scale, i.e., $m = 10, 20, 50, \dots$

MP and MOF as the fastest algorithms. The three algorithms have nearly the same complexity at higher dimensions (e.g., at $m = 1000$). We may then use *IDE-x* to achieve qualities near that of LP, while keeping the complexity as low as those of MOF and MP. Even with *IDE-s* the speed improvement is considerable.

D. experiment 4 - practical thresholds on sparsity

As stated in Section I, to ensure uniqueness of the sparsest solution, the number of active sources should be limited to $n/2$. But in practice, most methods breakdown before reaching this theoretical bound. In this experiment, we study practical limits (on the number of active sources) for *IDE-s*, *IDE-x* and LP.

In order to have more control over the sparsity, we generate source vectors according to a different model other than the Gaussian mixture. More specifically, given the number of active sources, $\#act$, a source vector is generated with exactly $\#act$ of its components randomly selected to be unity. The rest of the components, which represent inactive sources, are drawn from a zero-mean Gaussian with variance 0.01. This allows for a more accurate control of the sparsity. In fact, for this type of source, the quantity $\#act/(n/2)$ acts as a (normalized) measure of sparsity¹⁰ very useful to our discussion. Note that to ensure the ‘uniqueness of the sparsest solution’ property, $\#act/(n/2)$ should be kept below unity.

We will take $m = 1000$, $n = 400$ and select 25 values of $\#act/(n/2)$ in the range $[0.1, 1]$. For each $\#act$, both IDE’s and LP are applied to $N = 10$ realizations of the problem and the average SNR (over the N samples) obtained by each method is determined. Figure 6(a) illustrates the results when the general threshold sequence of experiment 2 has been used for both IDE’s.

Examining the figure, it is observed that the output SNR of both *IDE-s* and *IDE-x* is increased monotonically up to $\#act/(n/2) = 1/2$, after which it descends steeply¹¹ reaching nearly 0 dB around $\#act/(n/2) = 3/5$. The behavior of LP is somewhat similar except that the SNR begins to fall earlier and the degradation is more gradual. In particular, LP’s performance is still acceptable around $\#act/(n/2) = 3/5$. A general point to be made is that for all the three algorithms, *there seems to be thresholds on sparsity up to which they perform well and after which they degrade quickly in quality.*

It is possible to enhance the performance of IDE near the sparsity threshold by applying more iterations. To show this, we will examine the behavior using a longer threshold sequence with values spread wider across the $(0, 1)$ interval. The specific values are: $\epsilon = 0.9, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2, 0.1, 0.07, 0.05, 0.02, 0.01$. Figure 6(b) illustrates the results using this new threshold sequence. Note how IDE performance now degrades more gradually after $\#act/(n/2) = 1/2$, keeping the SNR at an acceptable level around $\#act/(n/2) = 3/5$; A behavior bearing more resemblance to LP.




Fig. 6. Average SSNR (in dB) versus normalized number of active sources, $\#act/(n/2)$, as a measure of sparsity. For each value of $\#act$, the average is obtained over $N = 10$ samples. The two plots correspond to different threshold sequences used in implementing IDEs: (a) 10-point sequence from experiment 2, (b) a more refined 13-point sequence.

Another interesting observation may be made by comparing the high-sparsity (i.e., low $\#act$) parts of the plots in Fig. 6(a) and (b): These parts are essentially unaffected by changing the threshold sequence. This result is in accordance with our previous intuitions. To sum up, *for relatively easy (i.e., highly sparse) problems, IDE is not sensitive to the choice of thresholds; Roughly general threshold sequences may be used without sacrificing performance; It is for difficult problems near the sparsity edge that the choice of threshold sequence really matters. In fact, the sparsity (edge) above which the method works is set by the chosen sequence.*

The observation we made that there is a threshold on $\#act$ (below the one suggested by theory) which limits the performance in practice has been pointed out by various authors. In fact, the figures we encountered for $\#act$ has also been obtained for the LP approach before. For example, [17] reported the experimental bound of $3n/10$ on $\#act$ for the minimum l^1 norm solution to coincide with the sparsest solution. The bound $n/4$ has been obtained for the incomplete Fourier dictionary in [2]. It appears that developing methods to fill the gap and work right up to the $n/2$ limit would be one of the challenges to be faced in the future.

E. experiment 5 - sensitivity to noise in the matrix

In SCA applications, where the \mathbf{A} matrix is estimated from mixture data, the robustness of the source-determination

¹⁰Again to be accurate, the quantity should be considered a measure of non-sparsity. To simplify discussion, however, we neglect these technicalities.

¹¹Some of the steepness is due to how the IDE’s have been implemented...




Fig. 7. Effect of noisy \mathbf{A} on performance: plots of average SSNR (in dB) versus average SNR of matrix \mathbf{A} . The averages are obtained over $N = 10$ noisy realizations of \mathbf{A} . The original \mathbf{A} is 200×500 .

algorithms to ‘estimation noise in \mathbf{A} ’ is important. This is not the case for applications like atomic decomposition where the dictionary \mathbf{A} is pre-determined. Even in these cases some noise may be induced on \mathbf{A} , for example, as a result of quantization. In this experiment, we will examine the effect of these perturbations on the performance of IDE’s and LP.

To model the perturbations, we will add to every component of the original matrix \mathbf{A} , a Gaussian noise of variance $\sigma_A \times \max |a_{ij}|$. The columns of \mathbf{A} are then re-normalized to unit l^2 norm¹². To conduct the experiment, we take a random source vector \mathbf{s} with $n/8$ of its components active (generated according to experiment 5 model), a random 500×200 matrix \mathbf{A} , and 10 values for σ_A in the interval $[0.001, 0.1]$. For each σ_A , we generate $N = 10$ noisy realizations $\{\hat{\mathbf{A}}_k(\sigma_A)\}_{k=1}^N$ according to the procedure mentioned above. An algorithm is then applied to the N noisy problems, designated with $\{(\mathbf{s}, \hat{\mathbf{A}}_k)\}$, resulting in the estimated source vectors $\{\hat{\mathbf{s}}_k(\sigma_A)\}_{k=1}^N$. Finally, the average (spatial) SNR in \mathbf{s} , i.e., $(1/N) \sum_{k=1}^N \|\mathbf{s}\|_2^2 / \|\mathbf{s} - \hat{\mathbf{s}}_k(\sigma_A)\|_2^2$, is plotted against the average SNR in \mathbf{A} , defined as $(1/N) \sum_{k=1}^N \|\mathbf{A}\|_F^2 / \|\mathbf{A} - \hat{\mathbf{A}}_k(\sigma_A)\|_F^2$ where $\|\cdot\|_F$ denotes the Frobenius matrix norm.

The results are illustrated in Fig. 7 for the algorithms IDE-s, IDE-x and LP. For IDE’s, the general sequence of experiment 2 has been used. A typical behavior is observed for the three algorithms: They resist small amounts of noise in \mathbf{A} (up to SNRs of nearly 30 dB), but they degrade quickly in quality as the noise is increased beyond some limit. Also note that *the quality gain of IDE’s over LP is only obtained for very low-noise \mathbf{A} matrices*. The SNR curves for the three algorithms converge as a result of an increase in \mathbf{A} -noise, indicating the loss of performance gain. Another notable observation is that, at high noise levels, IDE-x performs slightly better than both LP and IDE-s which is somehow suggestive of a ‘de-noising’ property. It may be attributed to the fact that IDE-x seeks to minimize the distance $\|\mathbf{x} - \hat{\mathbf{A}}\hat{\mathbf{s}}\|_2$ unlike IDE-s and LP which enforce $\mathbf{x} = \hat{\mathbf{A}}\hat{\mathbf{s}}$ on the solution; An equation that need not hold in the noisy cases.

VIII. CONCLUSION

We have shown that by (rough) detection of active sources, one can eliminate the need for a combinatorial search, effectively replacing it with one ‘comparison of an activity function against a threshold’ for each source. A possible choice for the activity function $g_i(\mathbf{x}, \hat{\mathbf{s}}^{(k)})$ was proposed based on ideas from binary hypothesis testing under Gaussian mixture prior for sources. The detection step required an estimate of the source vector, and together with an estimation step, it was used in an iterative setting to obtain the ‘Iterative Detection-Estimation’ family of algorithms. We proposed two approaches for source estimation (given that the sparsity pattern is roughly known): one was based on projection of the solution set of $\mathbf{x} = \mathbf{A}\mathbf{s}$ into the activity subspace in the ‘source space’ leading to the IDE-s algorithm. The other one was based on projection of \mathbf{x} on the subspace spanned by active atoms in the ‘mixture space’ which lead to the IDE-x algorithm.

We showed experimentally that with proper threshold selection, both versions of IDE can achieve accuracies comparable to LP (or even slightly better) after few iterations. The interesting point was that IDE’s achieve this much faster, with IDE-s (IDE-x) being nearly two (three) orders of magnitude faster than LP.

It was also observed that the algorithm is usually not ‘too sensitive’ to threshold values. In particular, a fixed threshold sequence may be used for every instance of a fixed problem family (determined by a fixed sparsity level and fixed n/m value), i.e., there is no need to modify the thresholds on a per-sample basis. Also, a threshold sequence was found experimentally that could be used over a wide range of problem families to produce ‘good’ results.

In general, these results suggest that IDE’s might be used as fast alternatives to LP when dealing with high-dimensional sparse decomposition problems. One might also think of IDE as a general framework of which the proposed algorithms are just two examples: There might be better ways of detecting (single) source activity, e.g. using better activity functions, thresholdless decisions (see below), etc. Similarly, there might be better implementations of the estimation step, e.g. using different cost functions.

For example, one may think about a thresholdless variant of IDE: we know from the uniqueness condition (Section I) that at most $n/2$ of sources may be active. Then, instead of using thresholds on the values of the activity function for detecting active sources, *all $n/2$ sources for which the values of the activity function are the highest are detected to be active*. Although using this approach no threshold is required, it makes the algorithm somehow ‘greedy’ (but of course not as greedy as MP). Consequently, the algorithm may get trapped in ‘local minima’, specially where the degree of sparsity decreases (this is verified by our first simulations). However, having no thresholds is advantageous enough to use such a version in some practical applications.

REFERENCES

- [1] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1999.

¹²The results were observed to be nearly the same without normalization.

- [2] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [3] D. L. Donoho and X. Huo, "Uncertainty Principles and Ideal Atomic Decomposition," *IEEE Trans. Inform. Theory*, vol. 47, no. 7, pp. 2845–2862, Nov. 2001.
- [4] D. L. Donoho and M. Elad, "Optimally sparse representation from overcomplete dictionaries via l^1 norm minimization," in *Proc. Nat. Acad. Sci.*, vol. 100, no. 5, 2003, pp. 2197–2202.
- [5] M. Elad and A. Bruckstein, "A generalized uncertainty principle and sparse representations in pairs of bases," *IEEE Trans. Inform. Theory*, vol. 48, no. 9, pp. 2558–2567, Sep. 2002.
- [6] J. J. Fuchs, "On sparse representations in arbitrary redundant bases," *IEEE Trans. Inform. Theory*, vol. 50, no. 6, pp. 1341–1344, 2004.
- [7] R. Gribonval and S. Lesage, "A survey of sparse component analysis for blind source separation: principles, perspectives, and new challenges," in *ESANN'06*, 2006.
- [8] R. Gribonval and M. Nielsen, "Sparse decompositions in unions of bases," *IEEE Trans. Inform. Theory*, vol. 49, no. 12, pp. 3320–3325, Dec. 2003.
- [9] Y. Li, A. Cichocki, and S. Amari, "Sparse component analysis for blind source separation with less sensors than sources," in *ICA2003*, 2003.
- [10] J. A. Tropp, "Greed is good: algorithmic results for sparse approximation," *IEEE Trans. Inform. Theory*, vol. 50, no. 10, pp. 2231–2242, 2004.
- [11] M. Zibulevsky and B. A. Pearlmutter, "Blind source separation by sparse decomposition in a signal dictionary," *Neural Computation*, vol. 13, no. 4, pp. 863–882, 2001.
- [12] M. Rudelson and R. Vershynin, "Geometric approach to errorcorrecting codes and reconstruction of signals," University of California, Davis, Department of Mathematics, 2005.
- [13] E. J. Candès and T. Tao, "Decoding via linear programming," *IEEE Trans. Inform. Theory*, to appear.
- [14] —, "Near optimal signal recovery from random projections and universal encoding strategies," California Institute of Technology, Applied and Computational Mathematics, 2004.
- [15] D. L. Donoho and J. Tanner, "Thresholds for the recovery of sparse solutions via L_1 minimization," in *Proceedings of the Conference on Information Sciences and Systems*, March 2006.
- [16] D. L. Donoho, "Compressed sensing," *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, April 2006.
- [17] —, "For most large underdetermined systems of linear equations the minimal l^1 -norm solution is also the sparsest solution," Tech. Rep., 2004. [Online]. Available: <http://www-stat.stanford.edu/~donoho/Reports/2004/>
- [18] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. on Signal Proc.*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [19] P. Bofill and M. Zibulevsky, "Underdetermined blind source separation using sparse representations," *Signal Processing*, vol. 81, pp. 2353–2362, 2001.
- [20] J.-F. Cardoso, "Blind signal separation: statistical principles," *Proceedings IEEE*, vol. 9, pp. 2009–2025, 1998.
- [21] J. Héroult and C. Jutten, "Space or time adaptive signal processing by neural networks models," in *Intern. Conf. on Neural Networks for Computing*, Snowbird (Utah, USA), 1986, pp. 206–211.
- [22] M. Babaie-Zadeh, C. Jutten, and A. Mansour, "Sparse ICA via cluster-wise PCA," *Neurocomputing*, vol. 69, pp. 1458–1466, August 2006.
- [23] M. V. Hulle, "Clustering approach to square and non-square blind source separation," in *IEEE Workshop on Neural Networks for Signal Processing (NNSP)*, Madison, Wisconsin, Aug 1999, pp. 315–323.
- [24] I. Daubechies, "Time-frequency localization operators: a geometric phase space approach," *IEEE Trans. Inform. Theory*, vol. 34, pp. 605–612, 1988.
- [25] L. L. Scharf, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Addison-Wesley, July 1991.
- [26] J. Nocedal and S. Wright, *Numerical Optimization*. New York: Springer, 1999.
- [27] N. Gould, M. Hribar, and J. Nocedal, "On the solution of equality constrained quadratic programming problems arising in optimization," *SIAM J. Sci. Computing*, vol. 23, no. 4, pp. 1375–1394, 2001.
- [28] A. A. Amini, M. Babie-Zadeh, and C. Jutten, "A new approach for sparse decomposition and sparse source separation," *EUSIPCO'06*, 2006.